

REMARKS/ARGUMENTS

In the Office action dated July 21, 2005, claims 1, 2, 5, 6, 8-26, 30-35 and 39-41 were rejected. Claims 3, 4, 7, 27-29 and 36-38 were objected to. Claims 1-41 remain pending in the application. Applicant respectfully requests further examination and reconsideration in view of the remarks set forth fully below.

I. Allowable Subject Matter

Applicant notes with appreciation that claims 3, 4, 7, 27-29 and 36-38 were objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims. Applicant has not rewritten these claims at this time in view of the below-provided remarks.

Applicant notes that although claims 14 and 18 are rejected, claim 14 includes the same subject matter as claim 3 and claim 18 includes the same subject matter as claim 7. Both claims 3 and 7 are objected as noted above. It appears as though claims 14 and 18 should have been included in the group of objected to claims.

II. Claim Rejections under 35 U.S.C. § 102

Within the Office Action, Claims 1, 2, 5, 6, 8-9, 11-20, 22-26, 30, 32-35, 39 and 40 were rejected under 35 U.S.C. 102(e) as being anticipated by U.S. Patent No. 6,567,417 to Kalkunte et al., hereinafter Kalkunte.

A. Independent Claim 1

Claim 1 of the present invention is directed to:

*"A method for managing packets in a shared memory buffer that serves a plurality of output links comprising:
writing packets into said shared memory buffer;*

linking said packets into a linked list in the order that said packets are written into said shared memory buffer;

using said linked list to identify the oldest packet that is intended for an output link that is available for packet transmission; and

dispatching, to the respective available output link, said identified oldest packet that is intended for *an available output link* even if said identified oldest packet is preceded by an older packet that is intended for *an unavailable output link*." (emphasis added)

Applicant asserts that claim 1 is not anticipated by Kalkunte. Support for this assertion includes an identification of relevant elements of Kalkunte followed by a comparison of claim 1 to Kalkunte.

Relevant Elements of Kalkunte:

- 1) The switch fabric (see Figs. 2 and 3) has eight port interfaces and each port interface has a 10 Gbps full-duplex connection to an external Gigabit-switch. (col. 4, lines 31 – 32, col. 4, lines 44 – 51)
- 2) Each of the eight port interfaces includes an ingress port and an egress port. The ingress ports enable packets to enter the switch fabric and the egress ports enable packets to exit the switch fabric. Each egress port includes a memory management unit (MMU). (see Fig. 3 and col. 5, lines 66 – col. 6, line 4)
- 3) The egress of packets from the switch fabric refers to packets exiting the switch fabric through one of the eight egress ports.
- 4) Fig. 10 is an expanded view of the MMU that is part of each egress port. Each MMU includes a Packet Pool (PP) which is "a 128 Kbyte SRAM used for the main storage of egress packets for that port." (col. 11, lines 22 – 23) The description provided with reference to Fig. 10 relates to the operation of one MMU at one of the egress ports. Each egress port has an MMU that operates independently of the other MMUs to manage the exit of packets through the respective egress port. (col. 8, lines 35 – 36 and 42 – 44) The egress port in Fig. 10 is identified for description purposes as

egress port 0, where the eight port interfaces and respective ingress/egress ports are identified as ingress/egress ports 0, 1, 2, 3, 4, 5, 6, and 7.

- 5) The egress of packets through an egress port is a 7-to-1 operation. That is, a particular egress port manages packets from the seven other ingress ports that need to exit via the particular egress port. In the example of Fig. 10, the MMU of egress port 0 manages the exit of packets received from ingress ports 1 – 7. This is shown in Fig. 10, where the MMU for egress port 0 receives packets from ingress ports 1 – 7. To reiterate, the egress of packets through a particular egress port is a 7-to-1 operation in which packets from seven ingress ports contend to exit from one egress port.

Comparison of Claim 1 to Kalkunte

Given the above-provided background, Applicant points out that the preamble of claim 1 recites “managing packets in a shared memory buffer that serves *a plurality of output links*.” As shown in Fig. 3 of Applicant’s specification, the subject matter of claim 1 is directed towards a situation in which a shared memory buffer buffers packets destined for one of multiple different output links. In contrast to claim 1, the MMU of an egress port as described with reference to Fig. 10 of Kalkunte is specific to one output link (i.e., egress port 0) and buffers packets from different input links (i.e., ingress ports 1 – 7) that are destined for the one output link (i.e., egress port 0). In the case of claim 1, when one output link is unavailable, the shared memory can still provide packets to other output links that may be available, while in Kalkunte if the one output link (i.e., egress port 0) is unavailable, then the shared memory (i.e., the Packet Pool (PP) from Fig. 10 of Kalkunte) is blocked from outputting packets. That is, the shared memory of claim 1 provides packets to multiple different output links while the shared memory of Kalkunte provides packets to a single output link. In this way, the operation addressed by claim 1 is fundamentally different from the operation addressed by Kalkunte.

Given the fundamental difference between claim 1 and Kalkunte, Applicant asserts that the distinction between an “available output link” and an

"unavailable output link" made in claim 1 is not anticipated by Kalkunte. In particular, because each egress port of Kalkunte operates independently to control the exit of packets through its one output link, the one output link is either available or unavailable. If the one output link is available, then the MMU determines the order in which buffered packets are dispatched. If the one output link is unavailable (i.e., because it is already being used to its maximum capacity), then no other packets can be dispatched. The possibility of simultaneously having an available output link and an unavailable output link does not exist from the perspective of an MMU in Kalkunte. Therefore, the step of "dispatching, to the respective *available output link*, said identified oldest packet that is intended for an available output link even if said identified oldest packet is preceded by an older packet that is intended for *an unavailable output link*" as recited in claim 1 makes no sense in the system of Kalkunte.

Further, Applicant asserts that Kalkunte does not disclose "linking *said packets* into a linked list in the order that said packets are written into said shared memory" as recited in claim 1. While Kalkunte does disclose linked lists, a close examination of Kalkunte reveals that in the switch fabric of Kalkunte, each packet is broken up into a group of cells and the cells of a single packet are linked together by a linked list. The concept of cells as being a portion of a single packet is disclosed by Kalkunte for example, at col. 8, lines 63 – col. 9, line 8, col. 15, lines 40 – 65, col. 18, lines 15 – 22, and Fig. 14. Linking the cells of a single packet using a linked list is disclosed by Kalkunte, for example, at col. 11, lines 1 – 21. In particular, "[e]ach queue entry points to the head of a packet in memory, and the rest of *the packet is maintained by a linked list* in the LLA." Further, at col. 11, lines 63 – 64, "[a] linked list is built (if not already done so) *for the packet*." Nowhere does Kalkunte disclose that multiple packets are linked together into a linked list. In sum, Kalkunte discloses a technique in which the cells that make up a single packet are linked by a linked list, whereas, claim 1 recites "linking packets into a linked list." (see col. 19, lines 1 – 28)

Further, the linked list of claim 1 maintains "the order that said packets are written into said shared memory buffer" and is used "to identify the oldest packet

that is intended for an output link that is available." In contrast, Kalkunte discloses using time stamps to identify the order of packets. In particular, at col. 11, lines 5 – 6, Kalkunte recites that the "XQ maintains a Time Stamp associating each packet in memory with an age value." The time stamp and not a linked list of packets is used by Kalkunte to identify the order of packets.

In sum, because Kalkunte does not disclose a shared memory that provides packets to multiple different output links or the linking of packets into a linked list, Applicant asserts that claim 1 is not anticipated by Kalkunte.

Dependent Claims 2, 5, 6, 8, 9, and 11 were also rejected under 35 U.S.C. § 102 as being anticipated by Kalkunte. Claims 2, 5, 6, 8, 9, and 11 depend from independent claim 1. As discussed above, Applicant asserts that independent claim 1 is allowable over Kalkunte. Claims 2, 5, 6, 8, 9, and 11 are therefore allowable as depending from an allowable base claim. Additional support for select claims dependent from Claim 1 is provided below.

Claim 2 is directed to:

"The method of claim 1 wherein linking said packets into a linked list includes establishing a NEXT pointer for each packet that identifies the next packet on the linked list."

As described above, Applicant asserts that Kalkunte does not disclose linking packets into a linked list. Because Kalkunte does not disclose linking packets into a linked list, it follows that Kalkunte does not disclose "establishing a NEXT pointer..." as recited in claim 2.

Claim 5 is directed to:

"The method of claim 1 further including establishing a linked list HEAD pointer that identifies the oldest packet that is written into said shared memory buffer."

As described above, Applicant asserts that Kalkunte does not disclose linking packets into a linked list. Because Kalkunte does not disclose linking packets into a linked list, it follows that Kalkunte does not disclose "establishing a HEAD

pointer..." as recited in claim 5. Further, as described above, Kalkunte discloses using time stamps to determine the relative age of packets, not a linked list.

Claim 8 is directed to:

"The method of claim 1 further including identifying the availability *of each of said output links*.

As described above, each egress port of Kalkunte has an MMU that independently manages the exit of packets through the respective egress port. The MMU of a respective egress port is only concerned with the availability of its one egress port and has no need to identify "the availability of each of said output links" as recited in claim 8.

Claim 9 is directed to:

"The method of claim 8 wherein packets are dispatched in dispatch cycles and further including assessing the availability *of each of said output links* at the beginning of each dispatch cycle."

As described above, each egress port of Kalkunte has an MMU that independently manages the exit of packets through the respective egress port. The MMU of a respective egress port is only concerned with the availability of its one egress port and has no need to identify the availability "of each of said output links" as recited in claim 9.

B. Independent Claims 12, 23, and 32

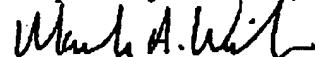
Claims 12, 23, and 32 include similar limitations to claim 1. Because of the similarities, Applicant asserts that the remarks provided above with reference to claim 1 apply also to claims 12, 23, and 32.

C. Dependent Claims 13 – 22, 24 – 31, and 33 – 41

Claims 13 – 22, 24 – 31, and 33 – 41 are dependent on claims 12, 23, and 32, respectively. Applicant asserts that claims 13 – 22, 24 – 31, and 33 – 41 are allowable based on the respective independent claims. Further, Applicant notes that the remarks made in reference to claim 2 apply also to claims 13, 25, and 34, the remarks made in reference to claim 5 apply also to claims 16, 26, and 35, the remarks made in reference to claim 8 apply also to claims 19, 23, and 39, and the remarks made in reference to claim 9 apply also to claims 20, 30, and 40.

Applicants respectfully request reconsideration of the claims in view of the foregoing remarks. A notice of allowance is earnestly solicited.

Respectfully submitted,



Mark A. Wilson
Reg. No. 43,994

Wilson & Ham
PMB: 348
2530 Berryessa Road
San Jose, CA 95132
Phone: (925) 249-1300
Fax: (925) 249-0111